

Automating UI Optimization through Multi-Agent Reasoning

Zhipeng Li

Department of Computer Science
ETH Zürich
Zürich, Switzerland
zhipeng.li@inf.ethz.ch

Christoph Gebhardt

Department of Computer Science
ETH Zürich
Zürich, Switzerland
christoph.gebhardt@inf.ethz.ch

Yi-Chi Liao

Department of Computer Science
ETH Zürich
Zürich, Switzerland
yichi.liao@inf.ethz.ch

Christian Holz

Department of Computer Science
ETH Zürich
Zürich, Switzerland
christian.holz@inf.ethz.ch

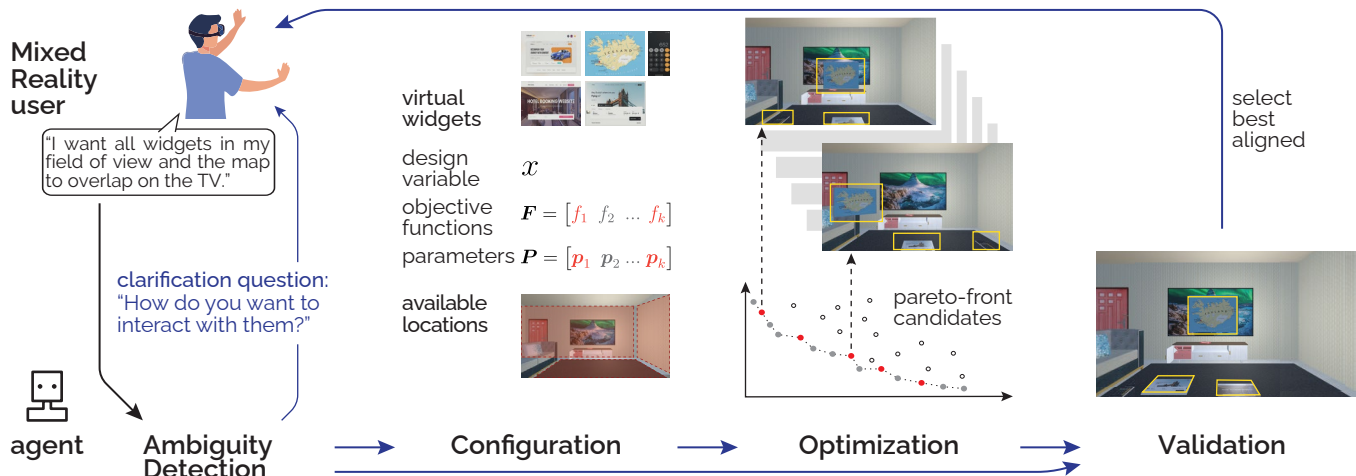


Figure 1: AutoOptimization is a framework for *personalizing optimization-based UI adaptations*. From the user’s stated instructions, we infer relevant terms and parameters to define a multi-objective optimization. First, *Ambiguity Detection* clarifies instructions with the user as needed, such that *Configuration* can specify the optimization problem by selecting the relevant objectives, UI elements, and placement locations. For *Optimization*, we run a solver to generate a set of Pareto-optimal layout candidates, and *Validation* selects the one layout from the set that best aligns with the user’s instructions.

Abstract

We present *AutoOptimization*, a novel multi-objective optimization framework for adapting user interfaces. From a user’s verbal preferences for changing a UI, our framework guides a prioritization-based Pareto frontier search over candidate layouts. It selects suitable objective functions for UI placement while simultaneously parameterizing them according to the user’s instructions to define the optimization problem. A solver then generates a series of optimal UI layouts, which our framework validates against the user’s instructions to adapt the UI with the final solution. Our approach thus overcomes the previous need for manual inspection of layouts and

the use of population averages for objective parameters. We integrate multiple agents sequentially within our framework, enabling the system to leverage their reasoning capabilities to interpret user preferences, configure the optimization problem, and validate optimization outcomes. We evaluate each step of our framework inside a Mixed Reality use case and demonstrate that *AutoOptimization* effectively increases the usability of UI adaptation schemes.

CCS Concepts

• **Human-centered computing** → **Mixed / augmented reality; Virtual reality; Interactive systems and tools.**

Keywords

Adaptive User Interfaces, Agent Systems, Mixed Reality.

ACM Reference Format:

Zhipeng Li, Christoph Gebhardt, Yi-Chi Liao, and Christian Holz. 2026. Automating UI Optimization through Multi-Agent Reasoning. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3772318.3791444>



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2278-3/2026/04

<https://doi.org/10.1145/3772318.3791444>

1 Introduction

Optimization has become a common approach for user interface (UI) adaptation. It operates by defining objective functions that specify what constitutes a good interface for a given application scenario, taking into account factors such as users' context or ergonomic considerations. A solver then adapts the UI, i.e., repositioning, showing, or hiding UI elements, to meet these objectives. Compared to manual exploration and adjustment, optimization has demonstrated a scalable and efficient solution for many UI adaptation problems, especially in Virtual Reality (VR) and Mixed Reality (MR) settings, where the interface design space is large [7, 8, 11, 13]. Recently, multi-objective optimization (MOO) has been introduced to tackle UI layout adaptation problems. Rather than identifying a single best solution, MOO generates a set of Pareto-optimal designs, each representing a unique balance between competing objectives such as visibility, accessibility, and contextual relevance [23]. These methods typically define fixed objective functions that incorporate scene semantics, physical affordances, or spatial consistency [7, 8, 11]. This enhancement broadens the applicability of optimization methods by supporting scenarios where multiple goals must be balanced.

Despite the advantages, conducting optimization remains a labor-intensive and complex process for both designers (or developers) and users. Designers must pre-define the design parameters (i.e., decision variables to be optimized) and the objective functions (i.e., metrics for evaluating design quality), all tailored to the scene and context. The objective terms and their parameters (such as usage frequency [30], arm length [8], or semantic association penalties [7]) are often hardcoded or based on population-level averages. Therefore, current optimization approaches cannot adapt to dynamic user needs or task-specific contexts, and instead rely on designers to configure them for each scenario.

For users, validating the final output can be equally demanding. When MOO is used, the result is often a collection of Pareto-optimal designs, each of which is a potential "best" design under a different trade-off. The user must manually evaluate and compare these options to determine which aligns best with their preferences [23]. While the weighted-sum approach can minimize the user's effort in exhaustively validating the Pareto-optimal designs [13], this approach relies on the assigned weights and potentially leads to a sub-optimal outcome. These steps show that optimization is not truly automatic. In fact, it often relies heavily on the designer's expertise (to set up the optimization task) and the user's effort (to explore and select from the output). Because of these manual burdens, optimization workflows cannot operate dynamically or autonomously across changing contexts without human supervision and reconfiguration. This brings us to a fundamental question: Can we *automate the optimization process to minimize the cumbersome aspects of UI adaptation for both designers and users*, and ultimately *enable a more dynamic, streamlined experience of adaptive UIs*?

In this paper, we introduce *AutoOptimization*, an optimization framework that leverages modern vision-language models (VLMs) as agents to streamline the cumbersome tasks of parameterizing and evaluating candidate designs in UI optimization. In this framework, the VLM plays the role of a *proxy designer*: First, it interprets the user's needs and the surrounding context, asking follow-up

questions when clarification is needed, and then automatically configures the optimization task. Second, when the solver produces a set of Pareto-optimal designs, the VLM evaluates these options and selects the final design based on its understanding of the user's intent. We call this framework "*AutoOptimization*" (short for Automating Optimization), as it is, to our knowledge, the first attempt to fully automate both the setup and decision-making phases of optimization through sequentially operating agents, reducing the user's involvement to the convenient act of supplying context verbally. Our goal is to shift the optimization workflow from being *human-driven* to *foundation model-assisted*, significantly reducing humans' effort in the process.

Our framework consists of two computational components: a vision-language model (VLM) and an optimization module. Both work in tandem to drive the UI adaptation process. The VLM agents work in a sequential, pipeline-based manner to interpret the user's verbal instructions, reason about the contextual intent, and make high-level design decisions. The optimization module then computationally generates candidate UI configurations. Finally, the VLM evaluates these results and selects the most suitable design.

Figure 1 illustrates the full process, starting with the user providing a high-level instruction or contextual goal, e.g., "I want all widgets in my field of view and the map to overlap on the TV." It then proceeds through four sequential steps: The first stage is *ambiguity detection*, where the VLM identifies potential ambiguities in the user input and seeks clarification through follow-up questions until it gathers enough information to configure the optimization task. Once the context is clarified, the VLM proceeds to *task initialization*, determining the relevant UI elements, design parameters, and objective functions based on the scene and user intent. The *optimization module* then solves the problem, producing a set of Pareto-optimal designs that represent trade-offs among multiple objectives. In the final step, the VLM *validates* these results and selects the most suitable design according to the original user-provided goals, effectively completing the end-to-end optimization process with minimal human intervention.

To evaluate the effectiveness of our *AutoOptimization* framework in a representative use case, we demonstrate its application to MR UI layout adaptation. To this end, we evaluated the individual stages of our framework. First, we collected data from experienced MR users on how they would instruct an intelligent assistant in various scenarios, categorizing their responses into well-informed or ambiguous instructions. By learning from the collected responses, the VLM can detect the ambiguity in instructions with an accuracy 91.26% in leave-one-user-out cross-validation and an accuracy of 92.93% in leave-one-scenario-out cross-validation. Next, we evaluated the performance of the VLM in selecting the layout that best aligns with user instructions from a set of Pareto-optimal candidates. We collected decisions from 26 participants and 26 instances of the VLM, evaluating 72 candidates across 18 scenarios. The results indicate that the VLM instances identified the most appropriate UI layouts similarly to the participants' selections. Finally, we evaluated our framework with an end-to-end user study. Participants customized MR UI layouts using *AutoOptimization*, a state-of-the-art baseline [24], and a purely manual placement method. The layouts generated by *AutoOptimization* demonstrated a closer alignment with user preferences, requiring

fewer and smaller adjustments compared to the baseline methods. Moreover, AutoOptimization achieved a comparable user satisfaction to manual placement while significantly reducing the effort. To summarize, we make these contributions in this paper:

- *AutoOptimization*, a novel framework that shifts UI adaptation from human-driven to foundation model-assisted workflows. Specifically, VLMs act as proxy-designers to automate key steps of an optimization process: input ambiguity detection, problem initialization, and result validation.
- an implementation of AutoOptimization for the use case of adaptive MR layout design. We also demonstrate its effectiveness and advantages over baseline approaches in a user study.

2 Related work

AutoOptimization is related to adaptive UIs in MR, their customization, and LLM use in HCI and problem solving.

2.1 Adaptive User Interfaces in Mixed Reality

Previous work has explored MR UI adaptation to the user's state [1, 14, 35, 44], for example by adjusting visibility, placement, or detail level based on cognitive load [33] or gaze activity [16]. Other work has optimized MR interfaces ergonomically using quantitative pose recommendations [12, 39].

Much MR interface optimization has focused on environment-based adaptation, aligning content with surroundings [9], either using geometry [21, 25] or specific objects like windows and walls [28]. SemanticAdapt adapted interfaces based on semantic relationships between virtual and physical objects [7], while InteractionAdapt considered input modes, positioning UI elements to leverage nearby surfaces for passive haptic feedback and facilitate cursor access [8]. Langerak et al. proposed a multi-agent reinforcement learning framework that dynamically adapts point-and-click user interfaces by coordinating specialized agents to optimize layout and interaction efficiency in real time [29]. More recently, MineXR was proposed to support researchers in designing MR interfaces across various contexts [10].

Although prior approaches effectively adapted layouts to the physical environment, they did not support tailoring UIs to individual users' needs and preferences, assuming instead that objectives and their parameters are fixed a priori across users and contexts. Because this assumption may not hold in real-world applications, our framework introduces the explicit capability of customizing the UI adaptation process based on the user's instructions, selecting a subset of objective terms and defining its parameters for optimal UI adaptation.

2.2 Customization of Adaptive User Interfaces

Adaptive MR systems often frame UI adaptation as a multi-objective problem, where static objective functions govern the placement of UI elements in the 3D environment. Since these functions apply uniformly to all users, prior research has focused on two approaches for accommodating user preferences: 1) adjusting objective weights and 2) exploring the Pareto frontier for optimal solutions.

1) A *weighted sum of objective terms* let users express the relative importance of each objective via weight adjustment. By manually

setting weights, the multi-objective problem is reduced to a single-objective one, yielding a single solution. For example, MenuOptimizer's Objective Space Panel allows designers to set weights for objectives like performance, consistency, and similarity [2], while OPTIMISM's GUI lets designers and end-users specify objective and heuristic weights for fabrication tasks [22]. In online 3D UI adaptation, pipelines such as SemanticAdapt [7], InteractionAdapt [8], and AUIT [13] provide Unity control panels for adjusting individual objective weights. Although this manual, trial-and-error process is straightforward, collapsing multiple objectives into one can obscure Pareto-optimal solutions.

2) *Pareto-optimal methods*, in contrast, generate the full set of solutions balancing potentially competing objectives, allowing users to explore and select preferred designs. For instance, ParetoAdapt visualizes the Pareto frontier and lets users choose designs that best match their preferences [24]. To reduce the manual effort of exploring trade-offs, our previous work proposed inferring user preferences from minimal UI element adjustments [43] and translated them into priority levels for a priority-based MOO algorithm that automatically selects the final design from the Pareto frontier. Other strategies include clustering to reduce solution sets, such as grouping haptic device designs by information transfer and recognition accuracy [32] or clustering input techniques by performance metrics [31]. Both approaches enable users to prioritize objectives before exploring solutions.

Unlike previous approaches, AutoOptimization eliminates the need for users to manually select from a set of designs. Instead, our framework builds on a dynamic formulation of the underlying MOO problem and guides candidate layout selection via the user's instructions. This way, AutoOptimization creates layouts that are tailored to the individual needs and preferences of users.

2.3 LLMs in HCI & Problem Solving

LLMs have been widely used in HCI research to simulate users and solve user-centric problems. In this paper, we leverage LLMs to reason about and simulate users' preferences to automatically solve UI optimization problems.

Before simulating preferences, LLMs must first gather signals to understand users' intentions. However, verbal instructions are often ambiguous. Therefore, LLMs need to detect and resolve ambiguities. Wang et al. demonstrated that LLMs can ask users for missing input and answer UI-specific questions without task-specific training, showing their capability to understand user instructions and UI layouts [45]. AmbigChat further decomposes ambiguous questions into a hierarchical "disambiguation tree," surfaces interactive UI widgets for each facet, and guides users through a conversational interface to clarify intentions [36].

Beyond ambiguity detection, LLMs can simulate user preferences. Aligning LLMs to human preferences has been explored in several contexts. Hämäläinen et al. prompted GPT to provide open-ended responses about video game experiences, finding results comparable to human participants [18]. Schmidt et al. noted that LLM-generated survey responses, while sometimes artificial, reveal insights into human preferences and survey design flaws [41]. Kang et al. showed that LLMs, when fine-tuned on small interaction datasets, can reach

or slightly surpass traditional recommendation models on predicting user ratings [26]. Similarly, Liu et al. integrated a crowdsourced library of user preferences into LLM reasoning and code generation, enabling automated UI widgets that align better with real user intentions compared to LLM-only methods [34].

LLMs are also used to tackle complex UI tasks. As one of the most classic methods, Chain-of-Thought prompting allows intermediate reasoning steps to guide problem-solving [48], and majority-voting across multiple chains further improves outcomes [46]. Following these, Tree-of-Thoughts extends LLMs' problem solving capability by structuring reasoning as a tree with self-evaluation, enabling backtracking and strategic decision-making [52]. As applications of these LLM-based reasoning frameworks, several studies investigate how to apply them to UI tasks. Guardian [40] uses LLMs to automate UI testing by prompting the model to plan each step needed to complete a task and then executing those steps to evaluate the interface. VisionTasker [42] uses a VLM to convert UI screenshots into natural-language descriptions, and then plan the next action, enabling end-to-end automated task execution.

In optimization-based UI adaptation systems, LLMs can be leveraged to directly tackle the optimization problem itself. Yang et al. demonstrated that LLMs can replace traditional optimizers by proposing candidate solutions in natural language and iteratively refining them, achieving competitive performance on small-scale optimization tasks [51]. Instead of fully replacing optimizers, another approach is to combine LLMs with existing optimization methods. For example, EvoPrompt integrates evolutionary algorithms with LLMs to optimize discrete prompts, leveraging both the search efficiency of evolutionary algorithms and the language coherence of LLMs [17]. Additionally, LLMs can provide parameters or contextual guidance for optimization tasks. For instance, SituationAdapt used a vision-language model to simulate user evaluations of virtual UI layouts [30], though it still relies on fixed weighted-sum objective functions, limiting its ability to adapt to dynamic user preferences.

Building on these prior works, our framework leverages a Chain-of-Thought approach to translate expressed user preferences into the formal context of a multi-objective optimization problem, automating the optimization-based UI adaptation process.

3 AutoOptimization framework

Our framework tightly integrates user-specific subjective instructions with a multi-objective optimization problem by means of an VLM-based agent to minimize human involvement. Figure 1 shows an overview of our framework, in which multiple agents are tasked with high-level decision-making to guide four phases:

For *Ambiguity Detection*, a first agent processes user instructions to extract essential information and preferences for adapting the UI; if it determines the user-provided input to be insufficient, it iteratively prompts the user for clarification and complete instructions. During *Configuration*, the second agent specifies the optimization problem, including involved UI widgets, prioritized objectives, and the corresponding parameter values. Our agent then runs an *Optimization* solver to search the Pareto front of the specified problem and generates multiple Pareto-optimal layout candidates. For *Validation*, a third agent validates these optimal layouts against the

user's original instructions and selects the most suitable layout. Below, we explain each of these components in more detail.

3.1 Ambiguity Detection: Clarifying user instructions for completeness

The input to our framework consists of an instruction through which a user describes how the UI should be structured, presented, or updated. We leverage verbal instructions, because humans are used to expressing their preferences in this manner, while VLMs can process such natural language as input for reasoning tasks.

When a user provides an instruction, the Ambiguity Detection agent evaluates whether the input is sufficiently clear and complete to infer layout preferences. To do so, the agent is also provided with the current UI layout, screenshots and names of all available widgets, and a screenshot of the user's field of view. This visual context enables the agent to ground the user's instruction in the actual interface state and to identify if essential details remain under-specified, including **which widgets are involved, how the user intends to interact with them, and user preferences**. The agent operates in an iterative refinement loop: if any ambiguity is detected, it asks the user a targeted clarification question, appends the response to the existing instruction, and re-evaluates the combined input. Through this accumulating context, the agent maintains a coherent understanding of the evolving instruction and continues the cycle until all relevant aspects are unambiguous, at which point it proceeds to the layout optimization stage.

The purpose of this dialog is to resolve ambiguity before proceeding to the actual optimization task. Our agent is equipped with ten clarification questions as an example for the VLM to respond to the user for resolving ambiguity in the original instruction. We adapted these questions from prior research [27, 37] and tailored them to the task of UI adaptation.

3.2 Configuration: Specifying the multi-objective optimization problem

With the complete set of user instructions to inform the layout adaptation task, our framework's second component translates them into a multi-objective optimization problem. For this purpose, a second agent dynamically includes or excludes individual terms from a set of optimization functions as well as UI elements, physical placement options, and other parameters based on the user's instructions. The resulting specification defines a customized optimization problem that affords a computational optimization for generating individualized, optimal UI solutions. Therefore, our agent overcomes a significant limitation of prior adaptive interface efforts, operating within the fixed optimization framework described in Sec. 1 & 2.2.

We define the general multi-objective optimization problem as:

$$\begin{aligned} \min_{\mathbf{x} \in X} \quad & \{f_1(\mathbf{x}, \mathbf{P}_1), f_2(\mathbf{x}, \mathbf{P}_2), \dots, f_k(\mathbf{x}, \mathbf{P}_k)\} \\ \text{subject to} \quad & g_j(\mathbf{x}) \leq 0, \quad j = 1, 2, \dots, m, \end{aligned} \quad (1)$$

where \mathbf{x} is the decision variable within the design space X , f_i is an objective function, and g_j denotes a constraint. The vector \mathbf{P}_i represents the parameters of the respective objective function. Therefore, our agent dynamically specifies the optimization problem during Configuration by including relevant objective functions $\{f_1, \dots, f_k\}$

and setting the values of the corresponding parameters P_i based on the user's instructions.

3.3 Optimization: Searching and generating Pareto-optimal layout solutions

At this stage of the framework, a computational solver is executed to solve the optimization problem specified in the Configuration phase. Our framework thus includes the parameterized subset of objective functions during Optimization and identifies the solutions to the specified MOO problem. The resulting Pareto-front consists of solutions that reflect various combinations of prioritized objectives. As a result, they represent a spectrum of optimal design candidates that are aligned with the user's expressed preferences.

Because the optimization may discover a large number of similar Pareto-optimal designs, a smaller set of representative designs to manage solution quantity and diversity is retained. Each of the designs we keep reflects a distinct trade-off within the objective space. We identify these designs using a scalarization-based filtering of the Pareto-optimal set (similar to Johns et al. [24]). We denote the final set of candidate solutions as X^* .

3.4 Validation: Selecting the one layout solution that best fits the user's instructions

In a final step, a third agent validates that the output of the optimizer aligns with the user's expressed preferences. Identifying the final UI layout x^* from X^* is challenging, because it requires assessing all solutions in X^* in the context of the user's (expressed) preferences.

We leverage the VLM-based agent's reasoning capabilities to evaluate the generated layouts in terms of the subjective instructions. Our agent compares all layout candidates in X^* against the user's initial instructions by prompting a custom VLM instance. This prompt comprises the user's instructions, an image of their physical surroundings, and the optimized virtual widgets. Our agent returns the layout $x^* \in X^*$ that best aligns with the instructions.

In summary, our agent-assisted selection from all Pareto-optimal results overcomes the previously needed manual effort [24] during UI adaptation. Our framework achieves this by validating layout candidates and picking the one that is most suitable UI following a user's prompt and in the context of their environment.

3.5 Aggregating user preferences over time

In addition to obtaining the final UI layout x^* , our framework also integrates user's expressed preferences over time. As a user provides new instructions, our framework combines this new input with all previous instructions. Aided by our Ambiguity Detection, this results in a comprehensive collection of instructions over time, which our agent leverages in every subsequent step. This allows our agent to interpret the user's instructions and desired adjustments with increasing accuracy, while ensuring that the optimized layouts resulting from our framework align increasingly better with the user's overall preferences.

4 Example use case: Mixed Reality UI layout adaptation

To evaluate the effectiveness of our *Auto-Optimization* framework in a representative use case, we apply it to the adaptation of user interface layouts in Mixed Reality. Unlike traditional interfaces confined to a single device, MR interfaces extend into the user's physical environment, enabling content to be positioned around the user in spatially meaningful ways. However, the vast design space of potential layouts, combined with constraints imposed by the physical environment and various contextual factors, makes automatic UI adaptation essential (as evident by numerous research efforts, e.g., [7, 8, 11]).

In this paper, we use this use case to concretely implement the previously conceptualized AutoOptimization framework, enabling us to evaluate its effectiveness and compare it to the current state of the art in UI optimization. In the following sections, we first describe the data collection process for fine-tuning the VLM-based agents, then present the use-case-specific implementation of our framework, and finally compare its performance to state-of-the-art approaches within the context of this use case.

5 Data Collection

To guide the adaptation of MR layouts based on verbal instructions, our approach requires relevant data. Since no such dataset existed, we conducted an online survey to gather input from MR users on how they would provide verbal instructions to an intelligent assistant that could layout their MR interface.

5.1 Survey Design

The survey was designed in alignment with the core concept of our framework: intelligent agents interacting with an MR UI multi-objective optimization scheme to enable easy customization of the layout adaptation process. For the optimization, we defined several objectives identified from related work that our framework should include: physical body strain and exertion [24], haptic feedback for touch interactions [8], UI design principles (such as alignment) [7], and the interface's relation to the physical world (including semantic relationships and overlay suitability) [7].

Before the survey, we explained to participants that their task was to provide instructions to an intelligent assistant capable of generating MR UIs based on those instructions. We explicitly outlined the objectives the assistant could consider and illustrated the process with two example instructions showing how the assistant could produce corresponding MR layouts. Participants then completed demographic questions before proceeding to the main survey.

Scenarios. Our survey featured 15 distinct scenarios in which participants were asked to provide instructions. Each scenario included a first-person perspective photo depicting the surroundings of a hypothetical MR user, along with a specific task and a set of virtual widgets. We designed three distinct sets of virtual widgets, each tailored to a specific context: work, entertainment, and travel planning. In each scenario, the task was designed to align with the selected widgets. Examples include scheduling a meeting (work widgets), online shopping (entertainment widgets), and booking a



Figure 2: Examples of scenarios from the data collection survey, showcasing virtual widgets and the simulated environment. Participants are asked to imagine themselves in the given background and to provide instructions to an intelligent assistant to arrange the virtual widgets that completes the specified task.

flight (travel planning widgets). The background photo was chosen to match the task setting, such as an office space, living room, or coffee shop. Figure 2 showcases three examples from our survey, illustrating the tasks and virtual widgets for each scenario.

Questions. In each scenario, participants were asked to provide instructions to the intelligent assistant on how to adapt the widgets to the background such that they could best perform the specified task ("Please provide the instructions to the Intelligent Assistant, so it can adapt the UI based on the current context and task.").

5.2 Participants

We recruited 29 participants (19 female, 10 male), ages 19–43 ($M=26$, $SD=3.42$) from an online crowd-sourcing platform. To guarantee a certain level of MR experience among participants, we screened them to ensure they used a MR device at least 1 to 5 times a month. Of those, participants reported their frequency of using MR: three participants reported daily use, six mentioned using it several times a week, eleven indicated they used it several times a month, and the remaining participants used it less frequently. Participants completed the survey in 30 min and received £4 as a gratuity.

We excluded participants who answered one or more of our three control questions incorrectly. Consequently, the data from 27 participants were used in the analysis.

5.3 Formative Analysis

From the 27 participants, we collected a total of 415 instructions, with an average length of 26.51 words ($SD=31.56$). The goal of our analysis was to determine whether these instructions could be directly utilized by our intended framework to guide a multi-objective UI optimization with the specified objectives. Three researchers independently coded all collected instructions. They categorized the instructions into two groups:

- (1) Well-formed instructions that included all necessary information, or had missing details that could be inferred from the context.

- (2) Incomplete or ambiguous instructions that lacked critical information which could not be inferred.

A majority vote among the three researchers was used to determine the final categorization of each instruction. As a result, 217 out of 415 responses were classified as well-formed, while the remainder were deemed ambiguous or incomplete. These results are consistent with findings from previous research [20, 37, 47] where participants also tended to provide incomplete instructions. This supports the validity of our design choice to incorporate an ambiguity detection module aimed at identifying ambiguous or incomplete input.

6 Use Case Implementation

In this section, we describe the implementation of the AutoOptimization framework for the specific use case of MR layout adaptation. While adhering to the general structure outlined in Section 3, we incorporate use case specific objective terms and constraints. Additionally, we leverage the data collected in Section 5 to adapt the VLM instances for this application.

6.1 Ambiguity Detection

We adapt the VLM used in this phase through few-shot learning by augmenting its context with data from our collection study. Specifically, we provide the module with all 405 pairs of collected instructions, along with their assigned categories (well-informed or ambiguous), to improve its understanding of the task. Appendix A includes the comprehensive prompt that accompanies this module.

After the Ambiguity Detection module resolves all ambiguous aspects of the user's instruction through clarification questions and the user's additional responses, it concatenates these user input into a single, complete instruction and passes it to the Configuration module in natural language. Keeping the instruction in natural language is an intentional design choice: it preserves readability and explainability, making the subsequent configuration decisions transparent and easier for human users to interpret.

6.2 Configuration

As described above, this agent specifies the optimization problem for which solutions will be identified in the next step of our framework. To tailor it to our specific use case, we incorporate well-established objective functions, their respective parameters, and constraints drawn from the adaptive MR layout literature. The customized agent is then employed to determine the relevance of these components based on the user’s verbal instructions. The following subsections provide a step-by-step explanation of this process.

6.2.a Objective functions

For the parameterizable optimization problem, we define the complete set of objective functions f_1, \dots, f_k using terms from prior work that have proven effective in generating usable MR layouts:

Spatial alignment [7] For each virtual widget, identify the voxels it occupies and compare them with the voxels occupied by other widgets. Calculate the distance between corresponding rows and columns of these voxels. A higher alignment cost indicates that the widget is poorly aligned with others, both vertically and horizontally. A lower cost signifies that the widget is well-aligned with neighboring widgets across both dimensions.

Field of view [7] For each virtual widget, calculate the angular difference between the widget’s position and the user’s gaze direction and multiply it with its observation probability. If the angular difference is less than 5 degrees (i.e., within the foveal region [38]), the cost is zero. A greater angular difference results in a higher cost, indicating that the widget is far from the user’s gaze direction, while a lower cost reflects near to the foveal area.

Anchor to physical object [13] If a virtual widget is anchored to a physical object, anchoring cost stems from the angular distance between the virtual widget and the physical object from the user’s perspective. A higher cost indicates that the virtual widget is more distant from its anchored object, whereas a lower cost suggests closer alignment. If the virtual widget is not anchored to any physical object, the anchoring cost is considered zero.

Overlay physical object [30] For each physical object, identify the voxels it occupies and check if any virtual widgets are positioned between these voxels and the user. If there is any, the overlaying cost is calculated as the angular differences between the virtual widget and the physical object, weighted by the object’s overlay suitability. A higher overlaying cost indicates that physical objects that are not well-suited for overlay are obscured by virtual widgets, while a lower cost suggests a clearer view of the physical objects.

Neck strain [24] For each virtual widget, calculate the angular distance between the user’s eye and the widget’s position on the x-z plane, and then multiply this distance by the widget’s observation probability. A higher cost indicates that widgets with higher observation probability are positioned above or below the user’s eye level, whereas a lower cost reflects widgets that are closer to the eye level.

Arm exertion [24] Similar to the neck strain objective, for each virtual widget, compute the angular distance between the

user’s shoulder and the widget’s position, and then multiply this distance by the widget’s interaction probability. A higher cost indicates that widgets with higher interaction probability are positioned significantly above or below the user’s shoulder level, whereas a lower cost reflects widgets that are closer to the shoulder level, thus enhancing ease of interaction.

6.2.b Fixed constraints

In addition to the objective functions, our optimization problem consists of fixed constraints $g_j(\mathbf{x})$ to ensure the optimized UI meets practical requirements. These include preventing occlusion, maintaining an appropriate field of view, and a reachable distance of elements. Specifically, these constraints include:

Occlusion [13] For each virtual widget, identify the voxels it occupies and check if any other virtual widgets are positioned between these voxels and the user. If there is any, the occlusion constraint cost is positive; otherwise, the cost is zero. This ensures that all widgets remain visible to the user.

Field of view [13] Similar to the field of view objective, each virtual widget must be within an angular distance of 60 degrees from the user to ensure it is within the user’s field of view. If the angular distance exceeds this threshold, the distance constraint cost is positive; otherwise, the cost is zero.

Distance [13] We assume that all virtual widgets should be within arm’s reach. To enforce this, we compute the distance between each virtual widget and the user’s shoulder, ensuring it does not exceed 0.65. If the distance exceeds this threshold, a positive distance constraint cost is applied; otherwise, the cost is zero. This threshold can be adjusted if direct touch interaction is not necessary.

6.2.c Parameterized objective functions

Our optimization problem includes the parameters of individual objective functions $\{P_1, \dots, P_k\}$ (P_i represents all parameters associated with objective f_i), relevant virtual widgets, and physical objects so as to better compute objectives and constraints. Specifically, these parameters include:

Interaction probability of a virtual widget [8] For each virtual widget, estimate the probability that the user will interact with it using direct touch.

Observation probability of a virtual widget [8] For each virtual widget, estimate the probability that the user will observe it.

Overlay suitability of a physical object [30] For each physical object, estimate the suitability of overlaying a virtual widget over it.

Anchored to physical object for each virtual widget [13] For each virtual widget, determine whether it is anchored to a physical object. If it is, identify the specific physical object to which it is anchored.

6.2.d Configuring the overall objective function

Based on the user’s instructions, our Configuration module specifies the optimization problem by selecting a subset of $\{f_1, \dots, f_k\}$ and setting the values of the corresponding parameters P_i . To this

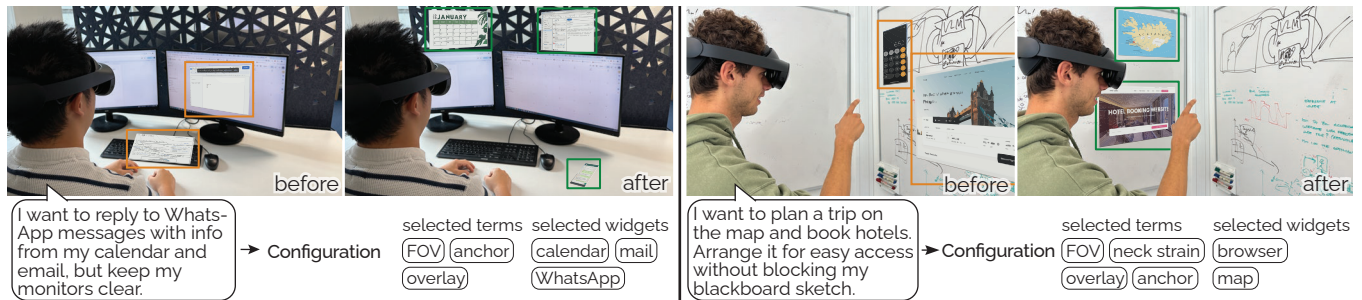


Figure 3: Illustrative demonstration of our Configuration component in a Mixed Reality context. Left: The user aims to reply to messages using information from virtual widgets while avoiding overlaying physical monitors. The Configuration module in our framework selects the objective terms that relate to the user’s expressed preferences, such that the Optimizer produces layouts that keep all widgets within the user’s field of view, the messenger anchored to the desk for haptic feedback, and the monitors visible and unoccluded. Right: The user plans a trip using a map and books a hotel while keeping the whiteboard sketch unobstructed. Our framework selects and displays the widgets relevant to the user’s instructions on the clean whiteboard.

end, we provide the respective agent with contextual information explaining the roles of objective functions, widgets, and their parameters. Based on this context, the agent first identifies the intended virtual widgets from a predefined set of candidates, thereby effectively defining a new design space Y with $\dim(Y) \leq \dim(X)$.

For example, in the left scenario illustrated in Figure 3, the Configuration module selects the "Calendar," "Mail," and "WhatsApp" widgets based on the user’s instructions, determining the set of virtual elements to be optimized. Similarly, based on the textual description of the available objective functions (f_1, \dots, f_k) that we provide, the module selects the subset of objectives $\{f_l(\mathbf{x}, \mathbf{P}_l), \dots, f_n(\mathbf{x}, \mathbf{P}_n)\}$ where $1 \leq l \leq n \leq k$, that best represent the adaptation goals expressed by the user. For instance, in the same scenario illustrated in Figure 3, the module selects the "FOV," "Anchor," and "Overlay" objectives based on the user’s instructions, ensuring that the optimization process prioritizes the most relevant adaptation criteria. Finally, we provide the module with an explanation of the associated parameters ($\mathbf{P}_1, \dots, \mathbf{P}_n$) of the selected subset of objective functions. Since selecting appropriate parameter values depends on the characteristics of both the virtual widgets and the physical objects, we also provide the agent with the relevant properties. The agent then estimates the values \mathbf{P}_i for $i = 1, \dots, n$ to best reflect the user’s input as captured by the previous module. In the same scenario shown in the left of Figure 3, the module determines a very low overlay suitability for the physical monitor, as the user prefers to keep their monitors clean. Similarly, in the right scenario, it selects the left blackboard as the anchor target for the virtual widget, since the user wants easy access to the widgets while avoiding obstruction of their sketch.

The Configuration module outputs its optimization settings in JSON format, which the Optimization module then reads and uses to configure the optimizer accordingly. Appendix B prints the full prompt that instructs the agent of our Configuration module.

6.3 Optimization and Validation

Neither the Optimization module nor the Validation module needs to be specifically tailored to a particular application. The Optimization module is responsible for solving the optimization problem that

was adapted to the specific use case by the Configuration module. It then solves the optimization problem, obtains a set of Pareto-front solutions, and sends them to the Validation module in JSON format. The Validation module receives the solution set and parses each candidate. It then applies each configuration to the UI layout and captures corresponding screenshots of the interface and physical context, all while keeping these intermediate states invisible to the user to avoid disrupting the interaction experience. It also receives the clarified user instruction to support subsequent evaluation. The Validation module then selects the most suitable solution by assessing both its alignment with the user’s instruction and its overall usability. Appendix C details the comprehensive prompt for the Validation module.

6.4 Technical Implementation

We implemented our system for the Meta Quest 3 using Unity 2021. The optimization module uses the AUIT framework [13], a toolkit designed to create adaptive Mixed-Reality applications. The toolkit integrates directly with Unity, using Unity GameObjects and their properties as inputs for optimization. We employ NSGA-III [5] to discover a set of optimal solutions on the Pareto-front based on the selected objectives (population size: 100, 40 generations). The solver is implemented in Python 3.9 using the pymoo [3] package and interacts with the Unity-based adaptation system. We apply the AASF decomposition[49] as the scalarization function to select a subset of solutions on the Pareto-front that balance different objectives effectively. The process is initialized with a set of reference vectors, computed via Riesz s -Energy[4], to ensure evenly distributed objective weight combinations. We control the size of the reduced candidate set by adjusting the number of reference vectors. We leveraged the GPT4 Vision 2024-05-01-Preview model of Azure OpenAI as VLM and access it via its Python API. The individual modules communicated via socket-based networking.

7 Evaluation

The objective of AutoOptimization is to streamline UI adaptation by automatically interpreting user instructions to define the optimization problem, specify inputs, and select suitable solutions. To




Figure 4: Classification accuracy of few-shot optimized VLMs with different number of participants and scenarios in training.

properly assess its effectiveness, we evaluate it in an MR layout adaptation use case. To this end, we first assess individual modules of the pipeline. Specifically, we examine whether our system can successfully detect ambiguities in user prompts and whether the comparison module selects optimal MR layouts that align with those chosen by MR users. Finally, we assess the overall usability by comparing our end-to-end pipeline with baseline adaptation methods in a user study.

7.1 Ambiguity detection

To assess if AutoOptimization can accurately detect ambiguities in user prompts, we utilized the data collected from our online survey to evaluate the performance of the ambiguity detection module.

Given that instructions could vary significantly across participants, we aimed to assess the scalability of few-shot learned VLMs, which serve as the underlying model of the Ambiguity Detection agent, when applied to new users. To achieve this, we performed a leave-1, 7, 13, 19, 26, 27-user-out cross-validation of the few-shot learned VLMs. As outlined in Section 6.1, the instructions and their corresponding categories from the training set were incorporated into the prompt. A VLM instance was then asked to categorize the instructions in the test set, and the classification results were compared to the researchers' decisions. Figure 4 (left) illustrates the results. Leaving all users (27) out resulted in an accuracy of 63.22% (SD = 7.20%), while including data from 7 users in the training set and leaving out 20 users resulted in an accuracy of 80.23% (SD = 7.40%), which significantly improved classification accuracy. With leave-one-user-out, the few-shot learned VLMs achieved a classification accuracy of 91.26% (SD = 4.69%). These results show that while the VLMs improve with additional user data, they already deliver strong performance using data from just half of the users.

Because users' instructions vary substantially across scenarios, we also investigate the capability of few-shot learned VLMs to detect ambiguity in previously unseen situations. With the instructions collected from 15 scenarios, we performed a leave-1, 4, 8, 11, 14, 15-scenarios-out cross-validation. We incorporated user prompts and their categories from the training scenarios and asked the VLMs to categorize instructions in the test scenarios, which simulated the detection of ambiguity in practical settings. The results are illustrated in the right panel of Figure 4. When all scenarios were excluded from the training set, the VLMs relied solely on their intrinsic capabilities for ambiguity detection, achieving an accuracy of 70.73% (SD = 11.28%). Incorporating data from 7 scenarios into the training set significantly improved accuracy to 88.09% (SD =

4.06%), and when only 1 scenario was excluded, the VLMs achieved an accuracy of 92.93% (SD = 2.05%). These results indicate that few-shot learned VLMs can effectively detect ambiguity in unseen scenarios by learning from a limited number of training scenarios.

7.2 Layout comparison

A key feature of our pipeline is its ability to enable textual guidance for the Pareto front exploration when searching for the optimal layout candidate. The assumption behind this feature is that AutoOptimization will evaluate the alignment between textual layout descriptions and layout candidates in a manner similar to human users. To assess this, we conducted an online survey to compare how experienced MR users and AutoOptimization rate the alignment of verbal descriptions and generated MR layouts.


7.2.1 Survey design. The survey was designed to evaluate how AutoOptimization and experienced MR users rate the alignment of textual layout descriptions with virtual UIs in different scenarios. Prior to the survey, participants were briefed on how AutoOptimization functions and how the MR layouts were generated. After this introduction, participants provided demographic information before proceeding to the main survey.

Scenarios. The main part of our survey included 18 scenarios. Each scenario consisted of four images of MR layouts rendered in the same environment, which were generated by AutoOptimization based on text prompts provided by participants of our data collection study. The text prompts were drawn from three specific environments: an office space, a living room, and an airplane.


To ensure consistency, we recreated the environments from our data collection study in Virtual Reality and used these replicas to render the MR layouts. Participants were instructed to treat the VR environments as real-world physical surrounding, with only the MR layouts were to be considered virtual.

The layouts represented Pareto-optimal design candidates generated by Pareto-AUIT based on input from our Configuration module. These selected solutions were rendered in VR, and screenshots were captured. Figure 5 illustrates the four candidate designs along with the corresponding text prompt for three sample scenarios, each rendered in one of the target environments.

Question. For each scenario, participants should select the layout that best aligned with the prompt from the four options provided



(a) Scenario 8 in an office generated using the prompt: "I'd like to schedule a meeting, open the calendar widget on the left and the email widget on the right of the laptop, and the messenger widget under the laptop, don't occlude the laptop." Design (a) is preferred by VLMS and PTPS.



(b) Scenario 4 in a living room generated using the prompt: "Place the map on the tv and put the flight booking and calculator on the desk so I can interact with it." Design (a) is preferred VLMS and PTPS.



(c) Scenario 16 in an airplane generated using the prompt: "I would like to watch videos on the tv and scroll and interact with the sports news, maybe play music at the same time, please place the spotify next to youtube." Design (d) is preferred by VLMS and PTPS.

Figure 5: Examples of scenarios from the layout comparison survey, each displaying four candidate solutions (a–d), the prompt it was generated with, and the design preferred by participants of the survey. The scenario numbers and candidate solutions (a–d) correspond to the respective scenarios and categorical values (A–D) displayed in Figure 6.

("Please select the layout that best aligns with the user's instructions."). The question was presented as alternative forced choice with four stimuli (4-AFC) [50].

7.2.2 Participants. We recruited 27 participants (15 female, 12 male), ages 18–54 ($M=35$, $SD=10.8$) from an online crowd-sourcing platform. To guarantee a certain level of VR experience among participants, we screened them to ensure they used a VR device at least 1–5 times a month. Of those, one participant reported using VR 11–15 times a month, 3 participants used it 6–10 times, and the remaining participants used it 1–5 times per month. Participants completed the survey in around 20 min and received £3 as a compensation.

We excluded one participant who answered two out of three control questions incorrectly. Consequently, the data from 26 participants (PTPS) were used in the analysis.

7.2.3 VLM ratings. We used the same scenarios, text prompts, and layout images to generate ratings with AutoOptimization. Each scenario's four images and corresponding textual prompt were fed into AutoOptimization, which returned its best layout (using the approach described in subsection 3.4). To ensure a balanced comparison, we generated results from 26 distinct VLM instances of the agent of the Validation module, matching the sample size of the human participants. For our analysis, this yielded a total of 468 ratings per condition (VLMS and PTPS).

7.2.4 Results. The goal of our analysis was to determine if AutoOptimization assesses alignment between verbal instruction and generated MR layout similar to the population of MR users. To this end, we computed the categorical distributions of the MR layout voted as best aligned by VLMS and PTPS in each scenario. As the categorical data is discrete, we analyzed the relationship between these distributions using the Chi-Squared Test. Our results suggested that there was no difference between the distributions of the ratings by VLMS and PTPS in any of the scenarios ($p > 0.05$).

In terms of the validity of our approach, an essential factor is whether the mode of the categorical distributions from both groups overlaps. We calculated the fraction of scenarios where the mode was consistent between VLMS and PTPS, and found that this was the case for all scenarios, yielding a perfect fraction of 1.0. A high degree of overlap and generally low divergence between the distributions can further be observed in Figure 6. Figure 5 shows three examples of the four candidate designs, highlighting the one preferred by both VLMS and PTPS. The scenario numbers and candidate solutions (A–D) correspond to the respective scenarios and categorical values (A–D) displayed in Figure 6.

7.3 User study

To evaluate the usability of AutoOptimization, we conducted a user study comparing it to two baseline methods: manual widget placement and choosing a layout from the Pareto front.

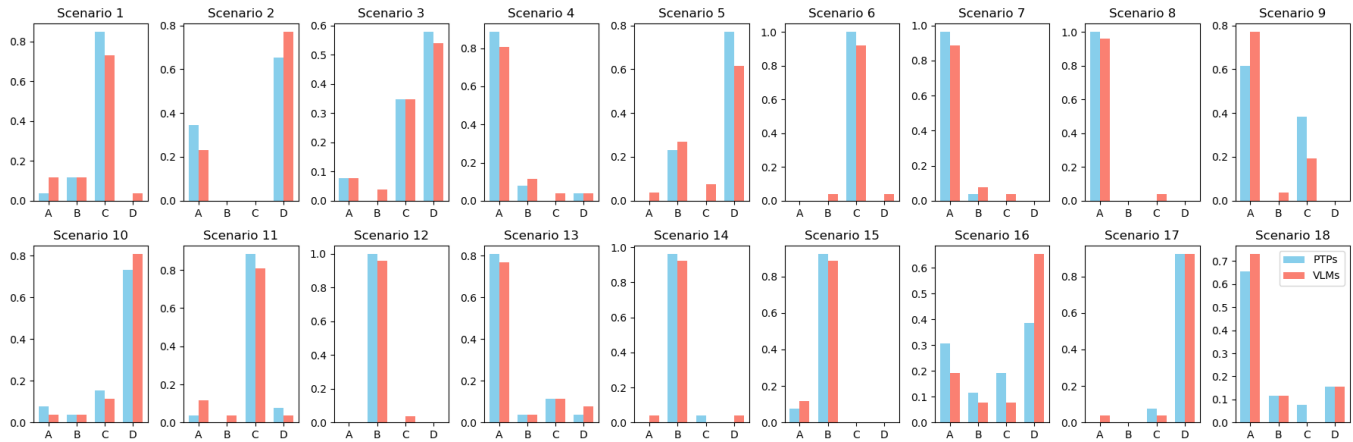


Figure 6: Categorical distributions of the MR layout voted as best aligned by vLMs and PTPs in each scenario.

7.3.1 Design. We employed a within-subject study design with one independent variable, *METHOD TYPE*, which had three levels: *AutoOptimization*, *ParetoAdapt*, and *ManualPlace*. To control for potential influence of environments, we introduced a control variable, *SCENARIO* (with three levels: living room, office, and airplane), representing the context in which participants were observed. In each condition, participants were asked to customize their user interface layout using each *METHOD TYPE* in all three *SCENARIOS*. They specified what they would like to do with the given widgets to the system first and then adjusted the generated user interface layouts. For dependent variables, we recorded the *number of adjustments* made to the user interface and the *adjustment distance* between the initial system-generated layout and the final layout refined by participants. We also employed the NASA-TLX questionnaire [19] to assess users' workload using each *METHOD TYPE*. Additionally, participants were asked to rank screenshots of the *ManualPlace*-adjusted interfaces, as well as the *AutoOptimization*- and *ParetoAdapt*-generated user interfaces (prior to any manual adjustment) at the end of each scenario. The order of both *SCENARIO* and *METHOD TYPE* was counterbalanced using a Latin square.

Scenarios. The same three virtual environments as in the layout comparison study were used: a living room, an office, and an airplane (as illustrated in Figure 5).

Methods. We compared *AutoOptimization* to both *ParetoAdapt* [24] and *ManualPlace*.

AutoOptimization Participants were asked to specify the task they aimed to complete with the virtual widgets in the current scenario and then provide instructions for the desired user interface layout. If the ambiguity detection module raised any clarification questions, participants responded accordingly. After our approach generated the MR layout, participants reviewed the interface and, if desired, refined it via manual drag-and-drop of widgets.

ParetoAdapt With *ParetoAdapt*, we modeled the optimization problem with pre-defined widgets, objective goals, and parameters. After the optimization finished, participants were presented with up to four candidate layouts from the Pareto front, selected using the AASF decomposition scalarization

function, and asked to choose the one that best aligned with their preferences. They had the option to refine the layout via widget drag-and-drop.

ManualPlace To begin, all virtual widgets were arranged in front of the user. Three widgets were placed at eye level, while the remaining two were positioned at hand level. This layout is comparable to the typical virtual environments on commercial platforms like Meta Quest and Apple Vision Pro. The order of the virtual widgets was randomized. Participants were given the option to modify the initial layout as needed. They could adjust the position and depth of the virtual widgets as needed and hide widgets by pressing a button on the controller.

7.3.2 Procedure. Participants started the study by providing informed consent and completing a demographic questionnaire. Following this, they engaged in a training trial designed to familiarize them with the available user interface elements and the experimental environment. During the training phase, participants were introduced to the workflow of all three *METHOD TYPES*. Subsequently, participants completed the conditions of the study, during which they provided instructions and observed the generated layouts across three scenarios for each method, resulting in a total of nine trials. After each trial, participants were asked to rank the layouts generated by each method. Naturally, the ranking of *ManualPlace* took place after the final layouts were created manually by the participants. For *AutoOptimization* and *ParetoAdapt*, the participants provided ranks before each final manual layout tuning; We chose to remove the factor of manual tweaking because our primary goal is to compare the quality of the layouts produced by these two methods against fully manual placement. After each condition, participants were asked to complete the NASA-TLX questionnaire. The entire study session was completed within 40 minutes.

7.3.3 Participants. We recruited 12 participants (4 female, 8 male), ages 22–28 ($M=25.91$, $SD=1.68$) from a local university. They reported their frequency of VR/AR headset use: four indicated using it several times a week; eight reported using it several times a month.

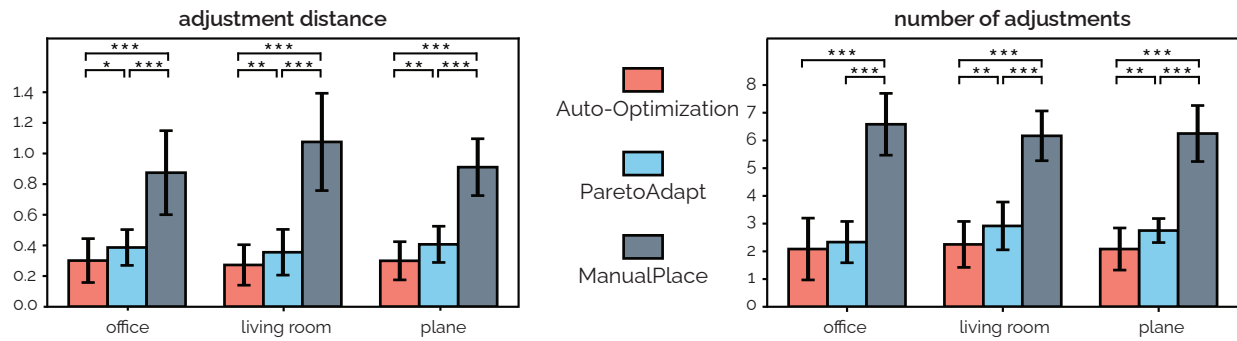


Figure 7: Means and standard deviations of the distance and the number of adjustments for each condition across all scenarios (*: $p < 0.05$, **: $p < 0.01$, *: $p < 0.001$).**

7.3.4 Apparatus. Participants wore a Meta Quest 3 headset, tethered to a desktop computer equipped with an Intel Core i7-12700K processor, an NVIDIA GeForce GTX 3070 GPU, and 32 GB of RAM. They interacted with the application using both a keyboard and the Meta Quest 3 controllers.

7.3.5 Results. We analyzed the effect of METHOD TYPE on three variables: *adjustment distance*, *number of adjustments*, and NASA-TLX scores, across all SCENARIOS. Since the assumptions of normality (Shapiro-Wilk $p < .05$) and sphericity (Levene’s $p < .05$) were violated for both *adjustment distance* and *number of adjustments*, we conducted a one-factor Aligned Rank Transform (ART) ANOVA for these two variables as well as to the NASA-TLX measures.

User interface adjustments. We found a main effect of METHOD TYPE on both *adjustment distance* [$F_{2,99} = 98.75, p < .001$] and *number of adjustments* [$F_{2,99} = 112.05, p < .001$] across SCENARIOS (Figure 7). Post-hoc tests revealed that participants made fewer adjustments [$t(11) = 2.60, p < .05$] and covered a shorter distance [$t(11) = 3.16, p < .01$] before reaching satisfaction with layouts generated by *AutoOptimization* compared to *ParetoAdapt*. For *ManualPlace*, participants began with a fixed central layout and manually adjusted all virtual elements, which could take multiple iterations. This process resulted in significantly higher adjustment counts and distances compared to both *AutoOptimization* [$t(11) = -13.44, p < .001$] and *ParetoAdapt* [$t(11) = -10.28, p < .001$]. Appendix D shows selected participant-generated UI layouts.

Layout preference. We observed a significant main effect of METHOD TYPE on participants’ preference rankings [$\chi^2(2) = 53.56, p < .001$] (Figure 8). Post-hoc test results indicated that *ManualPlace* ($M = 1.11, SD = 0.39$) outperformed both *AutoOptimization* ($M = 2.05, SD = 0.47$) [$Z = 64.5, p < .001$] and *ParetoAdapt* ($M = 2.83, SD = 0.37$) [$Z = 4.5, p < .001$], while *AutoOptimization* was ranked higher than *ParetoAdapt* [$Z = 87.5, p < .001$]. Participants always ranked the layouts generated through *ManualPlace* as their top choice. This is expected, as manual adjustment allowed them to directly position each widget to perfectly align with their implicit preferences. It is worth noting that layouts generated by *AutoOptimization* were consistently ranked higher than those generated by *ParetoAdapt*, highlighting that our approach generates layouts that are better aligned with the users’ preferences. Participants attributed their preference to the fact that, although *ParetoAdapt*

provided several reasonable options, these did not closely align with their ideal layouts. In contrast, *AutoOptimization* enabled participants to provide specific instructions, resulting in layouts that better matched their expectations.

Perceived workload. Figure 9 illustrates the results of the NASA-TLX. Our analysis revealed a main effect of METHOD TYPE on *Mental Demand* [$F_{2,33} = 3.46, p < .05$], *Physical Demand* [$F_{2,33} = 75.89, p < .001$], *Performance* [$F_{2,33} = 14.94, p < .001$], *Effort* [$F_{2,33} = 11.35, p < .001$], and *Overall workload* [$F_{2,33} = 8.47, p < .005$]. Post-hoc tests indicated that while participants perceived higher mental demand with *AutoOptimization* compared to *ManualPlace* [$t(11) = 2.88, p < .05$], they also reported lower physical demand [$t(11) = -12.30, p < .001$] and effort [$t(11) = -4.46, p < .001$]. Despite the increased mental demand, participants perceived no difference in performance using *AutoOptimization* and *ManualPlace* [$t(11) = 1.82, p > .05$]. Furthermore, *AutoOptimization* resulted in the lowest overall workload score compared to *ParetoAdapt* [$t(11) = -3.23, p < .01$] and *ManualPlace* [$t(11) = -3.82, p < .005$]. These results suggested that while *AutoOptimization* required more cognitive engagement, it alleviated the overall workload of customizing user interfaces while being perceived similarly effective.

Task completion time. We also analyzed the task completion time across METHOD TYPES. The results indicate that participants did not spend significantly more time using *AutoOptimization* ($M =$

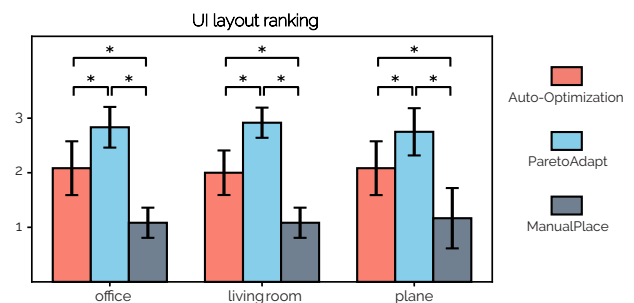


Figure 8: Ranking results of the layouts generated by the respective methods (lower is better). (*: $p < 0.05$, **: $p < 0.01$, *: $p < 0.001$).**




Figure 9: Means and standard deviations of the NASA-TLX questionnaire results for all three methods. (*: $p < 0.05$, **: $p < 0.01$, *: $p < 0.001$).**

46.83s, $SD = 12.03$) than *ParetoAdapt* ($M = 44.52s$, $SD = 5.38$) [$t(11) = -1.45$, $p > .05$], although completion times were notably longer compared to *ManualPlace* ($M = 11.47s$, $SD = 2.07$) [$t(11) = 13.05$, $p < .001$]. The additional time required for *AutoOptimization* can be attributed to participants needing to manually input instructions via a keyboard. This process could be significantly expedited through the integration of more efficient interaction modalities, such as speech-based input. In the case of *ParetoAdapt*, the extended completion time is primarily due to its optimization process, which requires substantially more time than that of *AutoOptimization*. This is because *ParetoAdapt* considers multiple objectives and explores a larger search space. However, despite the increased computational effort, it fails to produce results that better align with participants' customized needs, as evidenced by the greater number of UI adjustments required compared to *AutoOptimization*. Furthermore, NASA-TLX responses indicated that participants did not perceive a higher time demand when using *AutoOptimization* compared to *ParetoAdapt* and *ManualPlace*, suggesting that the system's time consumption was acceptable from a user experience perspective. Appendix E reports the runtime of each module of our system.

7.3.6 Ablation study: impact of disambiguating prompts. We conducted an ablation study to examine whether clarifying ambiguous prompts improved the quality of generated layouts. To this end, we generated intermediate results using participants' instructions before clarification, as illustrated in Figure 10.

In Figure 10(a)A, the participant stated, "I want to watch sports videos and also want to have some music." While there were two virtual widgets related to sports videos (Sky Sports and YouTube), the Configuration module selected only YouTube without ambiguity clarification. This decision conflicted with the participant's preference for Sky Sports, which better aligned with their intent to watch sports-specific content. By answering the clarification question, the user indicated the widgets they preferred, resulting in the layout shown in Figure 10(a)B.

Similarly, in Figure 10(b)C, the user did not specify a clear objective, stating, "I will interact with the calendar to scroll it up and down." Without ambiguity clarification, the system considered placing virtual widgets only on monitors and desks to provide haptic feedback for users. The clarification question prompted the user to provide more details about their objective. And when the user specified a preference for keeping the monitor and desk unoccupied, the optimizer placed the widgets in mid-air, which better accommodated the user's updated preferences.

Another type of ambiguity is illustrated in Figure 10(a)C, where the participant said, "I want to watch sports videos on YouTube and Sky Sports, so please place it on the monitor." The ambiguity arose from the user's use of "it," which left unclear whether they intended to place sports videos or YouTube on the monitor. This led the Configuration module to misinterpret the user's intent. After our ambiguity module prompted clarification, the user provided the necessary information, resulting in a layout that aligned with their preference.

8 Discussion

The goal of *AutoOptimization* is to enable users to customize their UIs using natural language by defining the optimization problem and selecting the optimal layout from a set of Pareto front solutions. In this section, we discuss the results of our evaluation, focusing on the framework's effectiveness in achieving this goal.

Detecting ambiguities in verbal commands. Our data collection study (Sec. 5) showed that users often give incomplete or ambiguous instructions. Our framework's Ambiguity Detection module demonstrated the capability to clarify them, achieving 91.26% accuracy in determining missing information in our leave-one-user-out cross-validation setting. We further tested the module with a leave-one-scenario-out cross-validation, where it achieved an accuracy of 92.93%. This result demonstrates its effectiveness in determining whether sufficient information has been gathered to generate a meaningful MR layout. The ablation study (Sec. 7.3.6) also showed

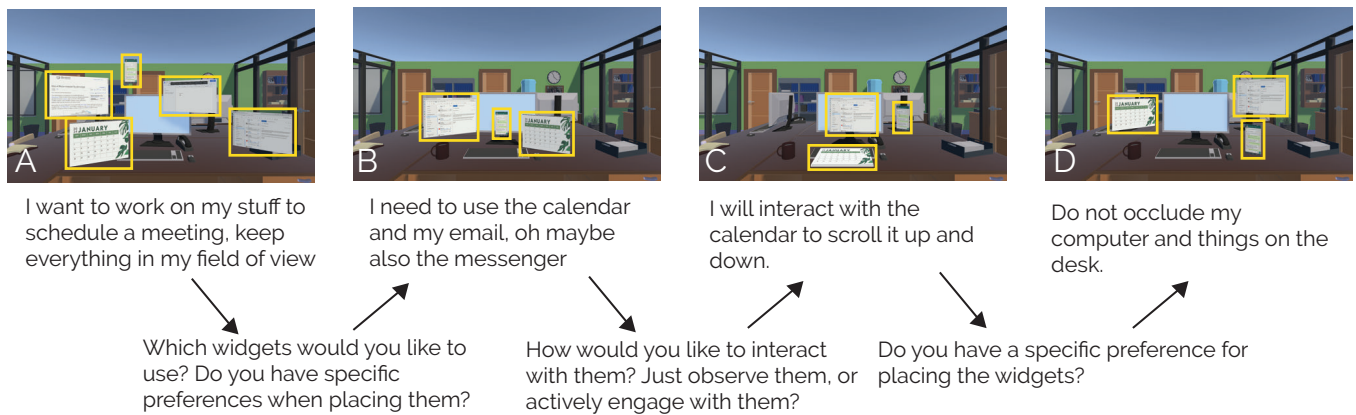
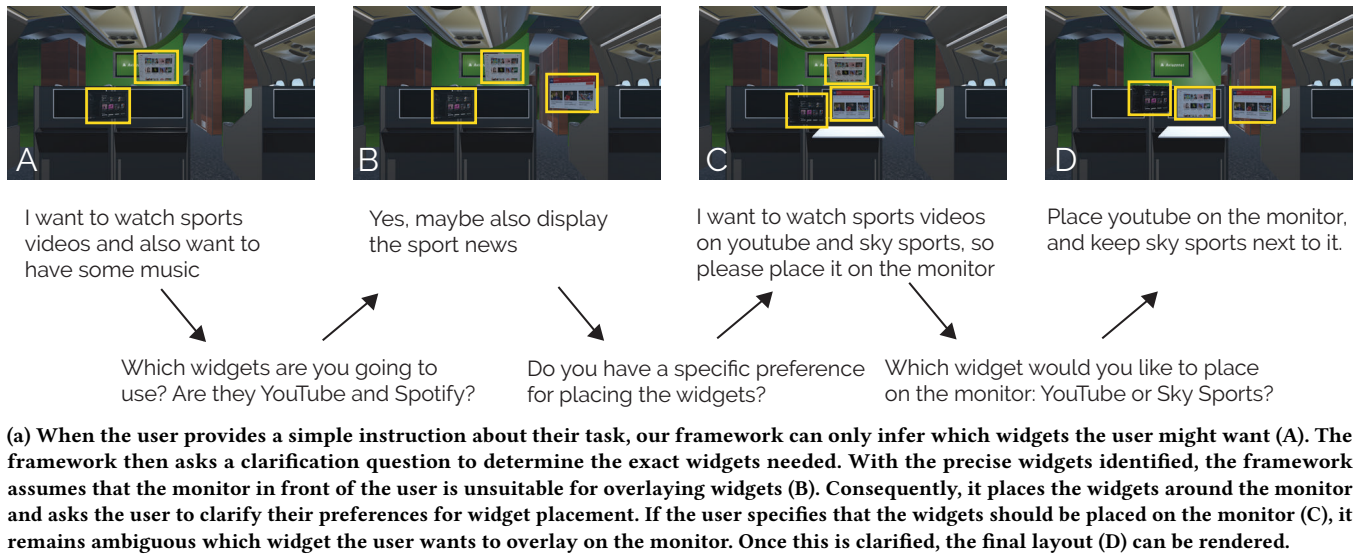


Figure 10: Examples of layouts generated from the user's instructions, which were partially ambiguous or incomplete. Following answers to our framework's clarification questions, the generated layouts were better aligned with the user's instructions (a–d). Note: Only the final layout (d) is presented to the user, whereas results (a–c) are only generated as examples.

that the UI layouts could be better tailored to users' preferences when they provided additional information in response to the clarification questions prompted by our Ambiguity Detection module.

users from a set of candidates, eliminating the need for them to manually evaluate and select layouts.

Validating layouts based on verbal instructions. In our layout comparison study, we assessed whether the Validation module accurately selects the final UI layout according to the user's instructions by comparing its selections with those made by experienced MR participants. The results show a 100% overlap between the layouts identified as best aligning with verbal instructions by both human participants and instances of Validation. This demonstrates that our framework successfully determines the most useful layout for

Customizing layouts with verbal instructions. Finally, we analyzed if our end-to-end pipeline could customize UIs based on users' verbal instructions more effectively than the current SOTA, *ParetoAdapt*, which relies on manual layout selection from the Pareto solution space of a static MOO problem. Results from our user study showed that participants made fewer adjustments and covered less distance with adjustments before reaching a satisfactory layout with *AutoOptimization* compared to *ParetoAdapt*.

Comparing to manual UI customization. In our user study, we compared *AutoOptimization* with a condition in which users manually adjusted the position of widgets with the controller. The NASA-TLX results revealed that while *AutoOptimization* led to a higher perceived mental demand, it significantly reduced both physical demand and general effort while achieving comparable performance to *ManualPlace*. Since *ManualPlace* allowed participants to adjust the UI layout precisely to their preferences, these findings suggest that *AutoOptimization* effectively enabled users to customize their UI layouts. We attribute the increased mental workload mainly to the effort of translating thoughts into verbal instructions for the system. To ease mental demand, future improvements could aim to reduce the number of iterations required to achieve users' goals by leveraging data from their past interactions or preferences.

In summary, study results suggest that our approach of using natural language to customize UI optimization helps users create layouts that better match their preferences while reducing the effort compared to baseline methods.

9 Limitation & Future Work

In the following, we discuss the limitations of our approach and potential directions for future research.

Objective space vs. design space instructions. We have demonstrated that *AutoOptimization* can generate UI layouts based on users' instructions. However, the framework is currently limited to instructions that relate to the objective space, meaning that instructions are interpreted only in terms of the objective functions currently supported by *AutoOptimization*. Commands such as "Place the video player three meters in front of me," which specify direct positions within the design space (i.e., the parameter space of each virtual widget's position), are not currently supported. To overcome this limitation, research should explore new objective functions that can map such design space instructions into the objective space. Another approach might involve preference optimization focused on learning and adjusting user preferences within the design space [50]. A hybrid approach, which infers whether the user's preferences relate to the objective space (e.g., minimizing neck strain) or the design space (e.g., placing elements in specific locations) and adapts the UI accordingly, would also be valuable.

Managing preferences before and after UI optimization. In our current implementation, *AutoOptimization* optimizes the UI layout based on users' interaction intentions, selected widgets, and personal preferences. These inputs are relatively high-level, typically describing general goals or preferences within the current UI layout. Using this information, *AutoOptimization* can reason about the user's preferences, configure the optimization problem accordingly, and select suitable solutions. However, we observed that even after optimization, users often want to make small manual adjustments, such as "move the calendar slightly to the left." These low-level adjustments are difficult to generalize across widgets or scenarios, as they often arise only after users have seen the layout results. This suggests that not all types of preferences are suitable for expression before optimization, which is also reflected in our study results. Accordingly, although *AutoOptimization* aims to interpret and reason about users' preferences for UI layout before optimization, it can

only disambiguate instructions and infer preferences to a limited extent. Certain preferences and requirements may only become apparent to users after observing the optimized UI, at which point further refinement is necessary.

Preferences elicitation for adaptive UI. A design feature in *AutoOptimization* that could influence users' preferences is the clarification conversation. The clarification questions are intended to resolve ambiguities in users' instructions, not simplifying the optimization problem, but to ensure that the optimizer is working towards the user's true intentions. However, when multiple candidate widgets or actions exist, clarification questions could subtly shape or elicit new preferences. In our current implementation, we designed the questions to avoid prompting users toward new preferences, but we acknowledge that they may still have encouraged consideration of aspects users had not previously thought about. For instance, asking whether the user wants to merely watch a widget or also interact with it could lead the user to refine or expand their interaction intention and preferences, which in turn might influence the optimization outcome. This highlights an inherent tension faced by adaptive systems: while clarification helps reduce ambiguity, it can also inadvertently shape or bias users' preferences. Future work should investigate methods to navigate this trade-off, ensuring that adaptive UI systems maintain both interpretive clarity and a faithful representation of users' genuine intentions.

User desired outcomes may not lie on Pareto front. In our current implementation, *AutoOptimization* employs six objective functions and three constraints, which are selected from prior work as effective signals for MR UI layout adaptation, but this fixed set may not generalize to all users, tasks, or contexts. Future research could extend or replace this set to better suit specific use cases. On the other hand, our optimization module identifies the Pareto front based on the existing objective functions and constraints, and the validation module selects a layout from this front that best aligns with the user's instructions. This workflow assumes that the user's preferred interface layout resides on the Pareto front, which may not always hold true. The core issue stems from the fact that the objective functions guiding the optimization process are merely mathematical approximations of user preferences. While these functions are designed to reflect user priorities, they inevitably fall short of capturing the full complexity of human judgment. Consequently, the Pareto front may not encompass all layouts that align with the user's true desires. To address this limitation, further development of objective functions and constraints that more accurately reflect user preferences could help narrow the gap between optimization outcomes and user expectations. Meanwhile, our approach supports manual post-optimization adjustments, allowing users to customize interface designs beyond algorithmic solutions.

Few-shot learning for Ambiguity Detection. Our ambiguity detection module uses few-shot prompting to help the VLM identify unclear aspects of users' instructions during UI layout optimization, which is in general challenging for VLMs [36]. However, prior work shows that in-context learning is highly sensitive to example selection [6, 15, 53], which can introduce biases that propagate into later configuration decisions. To mitigate this, we included all collected instructions and their annotations as examples, drawing

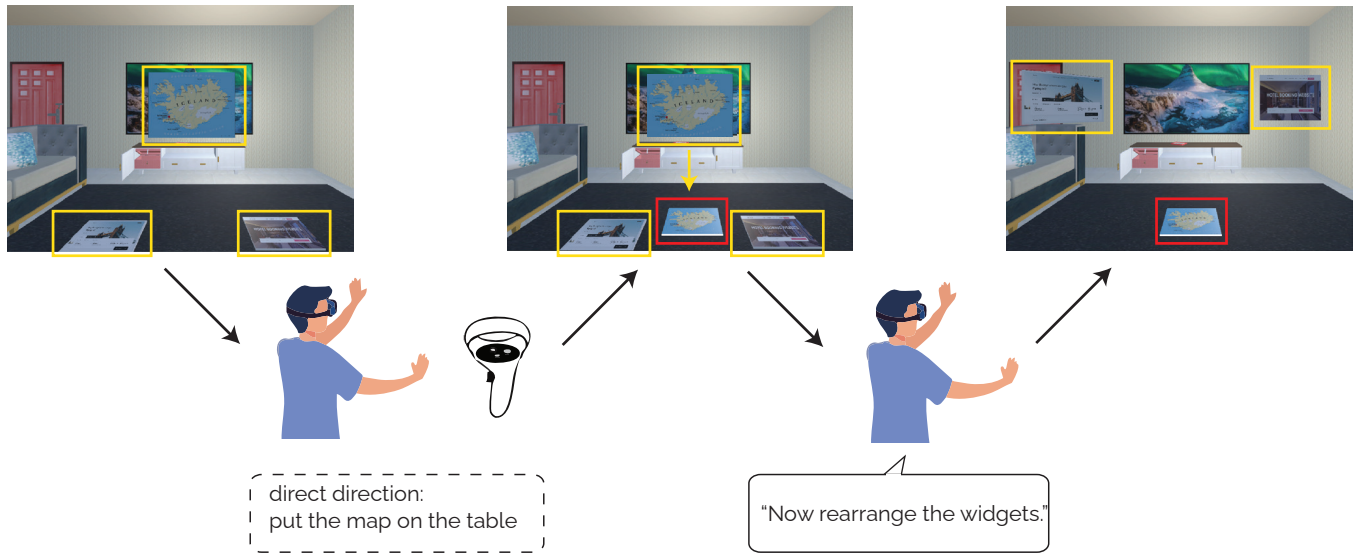


Figure 11: Illustration of integrating direct manipulation into AutoOptimization. When the user directly changes the position of an element (middle: moving the map from mid-air onto the table) and subsequently provides an additional instruction (e.g., “Now rearrange the widgets.”), AutoOptimization incorporates the user’s manipulation into the optimization process. The manipulation result is treated both as the starting point and as a strong constraint that must not be violated. For instance, moving the map away from the TV can be interpreted as a constraint to avoid blocking the TV. Based on this constraint, the system then optimizes the placement of the remaining elements (right).

from 29 participants and capturing a diverse range of phrasing styles. While this population-level prompt improves general robustness, it may still fail to capture individual user characteristics. A promising direction is to incorporate continual or user-specific few-shot learning: starting with population examples for initial interactions, and progressively augmenting the context with each user’s own past instructions. This approach could make ambiguity detection increasingly personalized over time, potentially improving accuracy and ensuring that subsequent optimization decisions are better aligned with individual preferences.

Evaluation of the Configuration module. Our framework comprises four components: Ambiguity Detection, Configuration, Optimization, and Validation modules. While the optimization module leverages well-established algorithms, we directly evaluated the ambiguity detection and validation modules through meaningful experiments. In contrast, evaluating the effectiveness of the Configuration module poses unique challenges. Based on user instructions, this module selects virtual widgets, defines the objective terms considered during optimization, and estimates their parameters.

A key challenge in evaluating the Configuration module lies in the absence of ground truth for objective term selection and parameter estimation. Since participants may not fully grasp the objective functions employed by the optimization module, they cannot serve as a reliable source of ground truth. Consequently, we assess the Configuration module only indirectly, through evaluation of the end-to-end system in our user study. The positive results observed, relative to baseline approaches, suggest that the

Configuration module effectively captures user intent and environmental context, generating UI configurations that align with users’ expectations.

Alternative approach to problem solving. Our framework utilizes multiple VLM instances (i.e., agents) within a Chain-of-Thought approach. However, the capabilities of current language models can be further enhanced by employing more complicated multi-instance structures [46, 52]. Specifically, we could restructure AutoOptimization as a tree-like hierarchy, where each module represents a node in the tree, generating multiple outputs based on its inputs.

Though computationally more demanding, this approach would enable the framework to self-evaluate, e.g., by comparing outcomes from different branches. This could result in two key benefits: improved alignment between layouts and user preferences, and reduced need for active user involvement. Future work could explore its impact on layout efficiency and accuracy.

Integrating direct manipulation into optimization. While AutoOptimization primarily relies on users’ verbal instructions to guide and optimize the UI layout, it also has the potential to incorporate additional forms of user input to further tailor the interface. For example, after customizing the layout based on verbal guidance, the user may interact with the interface to complete the current task, with the flexibility to adjust or reposition virtual elements during interaction. When a new round of UI optimization is initiated through verbal input, the framework does not restart from scratch. Instead, it incorporates the user’s direct manipulations and manual placements as strong constraints within the optimization process, adjusting the placement of remaining elements accordingly. This

ensures that both user-driven modifications and contextual considerations are preserved. Through this iterative process, the system progressively refines and adapts the interface, fostering a more coherent, personalized, and context-aware optimization workflow, as illustrated in Figure 11. We acknowledge the incorporation of additional forms of user input as important future work to further customize the automatic optimization process.

10 Conclusion

We have presented AutoOptimization, a framework that harnesses multi-agentic reasoning to enable personalized UI optimization through natural language instructions. By dynamically inferring relevant objectives and parameters from verbal instructions, our approach simplifies the complex task of preference-guided multi-objective layout optimization, reducing user effort while improving alignment with individual preferences.

Our evaluation in a MR use case confirmed the framework's effectiveness. The Ambiguity Detection module accurately flagged unclear instructions (91%), supporting reliable configuration and optimization downstream. The Validation module successfully selected final layouts that reflected participants' preferences, based on the outputs of our Optimization module. Together, these components ensured that the generated layouts closely matched user intent, requiring significantly fewer manual adjustments and yielding higher user satisfaction than the baseline method.

In tackling preference-based layout adaptation, AutoOptimization addresses limitations of prior approaches by optimizing UIs based on user preferences inferred from natural language instructions. By formulating the optimization problem according to these preferences and selecting the Pareto-optimal solution that best aligns with them, our system generates personalized UIs. This adaptive process reduces both task load and manual customization effort compared to prior approaches, underscoring the potential of AutoOptimization for creating individualized adaptive UIs.

Acknowledgments

Yi-Chi Liao was supported by the ETH Zurich Postdoctoral Fellowship Programme. Zhipeng Li was partially supported by the Swiss National Science Foundation (Grant No. 10004941).

References

- [1] Rawan Alghofaili, Michael S Solah, Haikun Huang, Yasuhito Sawahata, Marc Pomplun, and Lap-Fai Yu. 2019. Optimizing Visual Element Placement via Visual Attention Analysis. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 464–473. doi:10.1109/VR.2019.8797816
- [2] Gilles Bailly, Antti Oulasvirta, Timo Kötzing, and Sabrina Hoppe. 2013. Menuoptimizer: Interactive optimization of menu systems. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. 331–342. <https://doi.org/10.1145/2501988.2502024>
- [3] Julian Blank and Kalyanmoy Deb. 2020. Pymoo: Multi-objective optimization in python. *Ieee access* 8 (2020), 89497–89509.
- [4] Julian Blank, Kalyanmoy Deb, Yashesh Dhebar, Sunith Bandaru, and Haitham Seada. 2020. Generating well-spaced points on a unit simplex for evolutionary many-objective optimization. *IEEE Transactions on Evolutionary Computation* 25, 1 (2020), 48–60.
- [5] Julian Blank, Kalyanmoy Deb, and Proteek Chandan Roy. 2019. Investigating the normalization procedure of NSGA-III. In *International Conference on Evolutionary Multi-Criterion Optimization*. Springer, 229–240.
- [6] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [7] Yifei Cheng, Yukang Yan, Xin Yi, Yuanchun Shi, and David Lindlbauer. 2021. SemanticAdapt: Optimization-Based Adaptation of Mixed Reality Layouts Leveraging Virtual-Physical Semantic Connections. In *The 34th Annual ACM Symposium on User Interface Software and Technology (Virtual Event, USA) (UIST '21)*. Association for Computing Machinery, New York, NY, USA, 282–297. doi:10.1145/3472749.3474750
- [8] Yi Fei Cheng, Christoph Gebhardt, and Christian Holz. 2023. InteractionAdapt: Interaction-driven Workspace Adaptation for Situated Virtual Reality Environments. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–14.
- [9] Yi Fei Cheng, Tiffany Luong, Andreas Fender, Paul Strelci, and Christian Holz. 2022. ComforTable User Interfaces: Surfaces Reduce Input Error, Time, and Exertion for Tabletop and Mid-air User Interfaces. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 150–159.
- [10] Hyunsung Cho, Yukang Yan, Kashyap Todi, Mark Parent, Missie Smith, Tanya R Jonker, Hrvoje Benko, and David Lindlbauer. 2024. MineXR: Mining Personalized Extended Reality Interfaces. (2024).
- [11] Barrett Ens, Eyal Ofek, Neil Bruce, and Pourang Irani. 2015. Spatial constancy of surface-embedded layouts across multiple environments. In *Proceedings of the 3rd ACM Symposium on Spatial User Interaction*. 65–68.
- [12] João Marcelo Evangelista Belo, Anna Maria Feit, Tiare Feuchtnner, and Kaj Grønbaek. 2021. XRgonomics: Facilitating the Creation of Ergonomic 3D Interfaces. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems (Yokohama, Japan) (CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 290, 11 pages. doi:10.1145/3411764.3445349
- [13] João Marcelo Evangelista Belo, Mathias N Lystbæk, Anna Maria Feit, Ken Pfeuffer, Peter Kán, Antti Oulasvirta, and Kaj Grønbaek. 2022. Auit—the adaptive user interfaces toolkit for designing xr applications. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [14] Andreas Fender, David Lindlbauer, Philipp Herholz, Marc Alexa, and Jörg Müller. 2017. Heatspace: Automatic placement of displays by empirical analysis of user behavior. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. 611–621.
- [15] Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. Making pre-trained language models better few-shot learners. In *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing (volume 1: long papers)*. 3816–3830.
- [16] Christoph Gebhardt, Brian Hecox, Bas van Opheusden, Daniel Wigdor, James Hillis, Otmar Hilliges, and Hrvoje Benko. 2019. Learning Cooperative Personalized Policies from Gaze Data. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (New Orleans, LA, USA) (UIST '19)*. Association for Computing Machinery, New York, NY, USA, 197–208. doi:10.1145/3332165.3347933
- [17] Qingyan Guo, Rui Wang, Junliang Guo, Bei Li, Kaitao Song, Xu Tan, Guoqing Liu, Jiang Bian, and Yujiu Yang. 2023. Connecting large language models with evolutionary algorithms yields powerful prompt optimizers. *arXiv preprint arXiv:2309.08532* (2023).
- [18] Perttu Hämäläinen, Mikke Tavast, and Anton Kunnari. 2023. Evaluating large language models in generating synthetic hci research data: a case study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [19] SG Hart. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Human mental workload/Elsevier* (1988).
- [20] Jun Hatori, Yuta Kikuchi, Sosuke Kobayashi, Kuniyuki Takahashi, Yuta Tsuboi, Yuya Unno, Wilson Ko, and Jethro Tan. 2018. Interactively picking real-world objects with unconstrained spoken language instructions. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3774–3781.
- [21] Anuruddha Hettiarachchi and Daniel Wigdor. 2016. Annexing reality: Enabling opportunistic use of everyday objects as tangible proxies in augmented reality. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. 1957–1967.
- [22] Megan Hofmann, Nayha Auradkar, Jessica Birchfield, Jerry Cao, Autumn G Hughes, Gene SH Kim, Shriya Kurpad, Kathryn J Lum, Kelly Mack, Anisha Nilakantan, et al. 2023. OPTIMISM: Enabling Collaborative Implementation of Domain Specific Metaheuristic Optimization. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19. <https://doi.org/10.1145/3544548.3580904>
- [23] Christoph Albert Johns, João Marcelo Evangelista Belo, Clemens Nylandstedt Klokmoose, and Ken Pfeuffer. 2023. Pareto Optimal Layouts for Adaptive Mixed Reality. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems (Hamburg, Germany) (CHI EA '23)*. Association for Computing Machinery, New York, NY, USA, Article 223, 7 pages. doi:10.1145/3544549.3585732
- [24] Christoph Albert Johns, João Marcelo Evangelista Belo, Anna Maria Feit, Clemens Nylandstedt Klokmoose, and Ken Pfeuffer. 2023. Towards Flexible and Robust User Interface Adaptations With Multiple Objectives. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–17.
- [25] Brett Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira.

2014. Roomalive: Magical experiences enabled by scalable, adaptive projector-camera units. In *Proceedings of the 27th annual ACM symposium on User interface software and technology*. 637–644.
- [26] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. 2023. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv preprint arXiv:2305.06474* (2023).
- [27] Hyuhng Joon Kim, Youna Kim, Cheonbok Park, Junyeob Kim, Choonghyun Park, Kang Min Yoo, Sang-goo Lee, and Taek Kim. 2024. Aligning Language Models to Explicitly Handle Ambiguity. *arXiv preprint arXiv:2404.11972* (2024).
- [28] Wallace S Lages and Doug A Bowman. 2019. Walking with adaptive augmented reality workspaces: design and usage patterns. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*. 356–366.
- [29] Thomas Langerak, Sammy Christen, Mert Albaba, Christoph Gebhardt, Christian Holz, and Otmar Hilliges. 2024. Marlui: multi-agent reinforcement learning for adaptive point-and-click uis. *Proceedings of the ACM on Human-Computer Interaction* 8, EICS (2024), 1–27.
- [30] Zhipeng Li, Christoph Gebhardt, Yves Inglin, Nicolas Steck, Paul Strelci, and Christian Holz. 2024. SituationAdapt: Contextual UI Optimization in Mixed Reality with Situation Awareness via LLM Reasoning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–13.
- [31] Yi-Chi Liao. 2021. Computational Workflows for Designing Input Devices. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–6. <https://doi.org/10.1145/3411763.3443428>
- [32] Yi-Chi Liao, John J Dudley, George B Mo, Chun-Lien Cheng, Liwei Chan, Antti Oulasvirta, and Per Ola Kristensson. 2023. Interaction Design With Multi-objective Bayesian Optimization. *IEEE Pervasive Computing* 22, 1 (2023), 29–38. <https://doi.org/10.1109/MPRV.2022.3230597>
- [33] David Lindlbauer, Anna Maria Feit, and Otmar Hilliges. 2019. Context-aware online adaptation of mixed reality interfaces. In *Proceedings of the 32nd annual ACM symposium on user interface software and technology*. 147–160.
- [34] Yimeng Liu, Misha Sra, and Chang Xiao. 2024. CrowdGenUI: Aligning LLM-Based UI Generation with Crowdsourced User Preferences. *arXiv preprint arXiv:2411.03477* (2024).
- [35] Feiyu Lu and Yan Xu. 2022. Exploring spatial UI transition mechanisms with head-worn augmented reality. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [36] Jiaju Ma, Lei Shi, Kenneth Aleksander Robertsen, and Peggy Chi. 2025. AmbigChat: Interactive Hierarchical Clarification for Ambiguous Open-Domain Question Answering. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–18.
- [37] Ninareh Mehrabi, Palash Goyal, Apurv Verma, Jwala Dhamala, Varun Kumar, Qian Hu, Kai-Wei Chang, Richard Zemel, Aram Galstyan, and Rahul Gupta. 2023. Resolving ambiguities in text-to-image generative models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 14367–14388.
- [38] Michel Millodot. 2014. *Dictionary of optometry and visual science E-book*. Elsevier Health Sciences.
- [39] Roberto A. Montano Murillo, Sriram Subramanian, and Diego Martinez Plasencia. 2017. Erg-O: Ergonomic Optimization of Immersive Virtual Environments. In *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology (Québec City, QC, Canada) (UIST '17)*. Association for Computing Machinery, New York, NY, USA, 759–771. doi:10.1145/3126594.3126605
- [40] Dezhi Ran, Hao Wang, Zihe Song, Mengzhou Wu, Yuan Cao, Ying Zhang, Wei Yang, and Tao Xie. 2024. Guardian: A runtime framework for LLM-based UI exploration. In *Proceedings of the 33rd ACM SIGSOFT International Symposium on Software Testing and Analysis*. 958–970.
- [41] Albrecht Schmidt, Passant Elagroudy, Fiona Draxler, Frauke Kreuter, and Robin Welsch. 2024. Simulating the Human in HCD with ChatGPT: Redesigning Interaction Design with AI. *Interactions* 31, 1 (2024), 24–31.
- [42] Yunpeng Song, Yiheng Bian, Yongtao Tang, Guiyu Ma, and Zhongmin Cai. 2024. Visiontasker: Mobile task automation using vision based ui understanding and llm task planning. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. 1–17.
- [43] Yao Song, Christoph Gebhardt, Yi-Chi Liao, and Christian Holz. 2025. Preference-Guided Multi-Objective UI Adaptation. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. Association for Computing Machinery, New York, NY, USA, Article 120, 13 pages. doi:10.1145/3746059.3747645
- [44] Rafael Veras, Gaganpreet Singh, Farzin Farhadi-Niaki, Ritesh Udhani, Parth Pradeep Patekar, Wei Zhou, Pourang Irani, and Wei Li. 2021. Elbow-anchored interaction: Designing restful mid-air input. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [45] Bryan Wang, Gang Li, and Yang Li. 2023. Enabling conversational interaction with mobile ui using large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [46] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171* (2022).
- [47] Dulanga Weerakoon, Vigneshwaran Subbaraju, Nipuni Karumpullii, Tuan Tran, Qianli Xu, U-Xuan Tan, Joo Hwee Lim, and Archan Misra. 2020. Gesture enhanced comprehension of ambiguous human-to-robot instructions. In *Proceedings of the 2020 International Conference on Multimodal Interaction*. 251–259.
- [48] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
- [49] Andrzej P Wierzbicki. 1982. A mathematical basis for satisficing decision making. *Mathematical modelling* 3, 5 (1982), 391–405.
- [50] Kenta Yamamoto, Yuki Koyama, and Yoichi Ochiai. 2022. Photographic lighting design with photographer-in-the-loop bayesian optimization. In *Proceedings of the 35th annual ACM symposium on user interface software and technology*. 1–11.
- [51] Chengrun Yang, Xuezhi Wang, Yifeng Lu, Hanxiao Liu, Quoc V Le, Denny Zhou, and Xinyun Chen. 2023. Large language models as optimizers. In *The Twelfth International Conference on Learning Representations*.
- [52] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems* 36 (2024).
- [53] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*. PMLR, 12697–12706.

A Prompt to support Ambiguity Detection

We defined the following prompt to establish the context for the Vision-and-Language Model (VLM). The purpose of this prompt is to detect ambiguity in the user's command and, if needed, respond to the user's command with a follow-up question for additional information.

```

1 To evaluate whether the user's instructions provide sufficient information
  for customizing Mixed Reality user interfaces, focus on the following
  criteria:
2
3 Primary Widgets: The instructions should clearly identify the virtual
  widgets that users want to interact with, based on the task at hand. For
  instance, a user might specify they need a virtual button, slider, or menu.
4 Interaction Intention: The instructions should describe how users intend to
  interact with these widgets. This includes whether they want to use direct
  touch, voice commands, or simply observe the widgets.
5 User Preference: The instructions should reflect the user's considerations
  and priorities related to the placement and configuration of virtual
  widgets. Preferences might include:
6 Strain and Exertion: Minimizing physical effort or strain.
7 Field of View: Ensuring widgets are within an optimal viewing area.
8 Widget Alignment: Aligning widgets in a way that feels natural and
  intuitive.
9 Haptic Feedback: Incorporating feedback mechanisms if necessary.
10 Semantic Connection: Ensuring widgets are logically connected to their
  functions.
11 Surface Alignment: Placing widgets on appropriate surfaces.
12 Situation Context: Adapting widgets to the context in which they will be
  used.
13
14 Categories of Instructions:
15 Instructions with Sufficient Information: These include clear details on
  primary widgets, interaction intention, and user preferences. For example,
  "I need a virtual control panel with sliders for adjusting volume and
  brightness, which should be interactive through direct touch. The panel
  should be aligned to the bottom-right of the field of view, considering
  minimal strain."
16 Instructions with Missing, Ambiguous, or Irrelevant Information: These lack
  clarity or completeness in one or more of the key areas. For example, "I
  want an interface that works well." This instruction does not specify what
  widgets are needed, how they should be interacted with, or any preferences
  for their placement or configuration.
17
18 Task: Evaluate the following instructions and determine whether they are
  sufficient or fall into the category of missing, ambiguous, or irrelevant
  information based on the criteria provided.
19
20 The user's instructions are: [user's instructions]
```

We created the following prompt to provide the VLM with the user's instructions and its corresponding category.

```

1 instructions with information about the primary virtual widgets,
  interaction intention and user preference: \dots
2 instructions with missing, ambiguous or irrelevant information: \dots
```

B Prompt to support Configuration

Our framework input the following prompt into the VLM alongside the user's instructions for the purpose of selecting a subset of objectives.

```

1 The current user interface widgets include: [list of widgets]. Based on the
  user's instructions, please assess whether the following objectives should
  be considered in the user interface design. For each objective that should
  be considered, estimate its parameters as described:
2
3 Field of View Objective: Determines if virtual widgets should be positioned
  within the center of the user's field of view. The interval parameter,
  specified in degrees, represents the user's foveal range.
4
5 Arm Angle Objective: Indicates whether virtual widgets should be placed
  closer to the user's hand position, which is typically lower than the eye
  position, to facilitate easier interaction.
6
7 Neck Angle Objective: Determines if virtual widgets should be positioned
  near the user's eye level to improve visibility.
8
9 Alignment Objective: Assesses whether virtual widgets should be aligned
  with each other to create a structured user interface. The XTolerance and
  YTolerance parameters, in meters, define the allowable positional
  deviations on the X and Y axes.
10
```

```

11 Physical Anchor Objective: Evaluates whether virtual widgets need to be
  aligned with a specific physical area or object based on the user's
  instructions.
12
13 Overlay Objective: Determines if virtual widgets should obscure physical
  objects.
14
15 Please provide a detailed analysis based on the description of the user's
  instruction provided below: [User's instruction description]"]
```

We created the following prompt for the VLM together with the user's instructions to select virtual widgets and estimate parameters for objective terms.

```

1 The current user interface includes [number of widgets] virtual widgets.
  The widgets are: [list of widgets]. The physical areas involved are: [list
  of areas]. Based on the user's instructions, please determine the following
  parameters for each virtual widget:
2 Interaction Frequency: Represents the user's potential interaction
  intention with the virtual widget, ranging from 0 (no intention) to 1 (high
  intention).
3 Observation Frequency: Indicates the user's potential observation intention
  with the virtual widget, also ranging from 0 to 1.
4 Physical Anchor: Specifies whether the virtual widget is anchored to a
  physical area. The anchor must correspond to one of the provided physical
  areas. A virtual widget can have no anchor or be assigned to one physical
  area, and a physical area can be assigned to no more than one virtual
  widget.
5 Enabling Setting: Indicates whether the virtual widget will be used based
  on the user's instructions. If false, the widget will be disabled (i.e.,
  not displayed); if true, the widget will be enabled and visible.
6
7 For each physical area, assess the suitability of the area to be occluded
  or overlaid by virtual widgets based on the user's instructions. If a
  physical area is unsuitable for occlusion or overlaying, it cannot serve as
  a physical anchor for any virtual widget. The default overlay suitability
  should be set to 1 unless specified otherwise.
8
9 Please provide the analysis in the same order as the widgets: [list of
  widgets].
10
11 The user's instructions are: [user's instructions]
```

C Prompt to support Validation

Our framework prompts the VLM with the following query as well as the user's instructions to validate candidate outcomes from the optimizer on the Pareto-front.

```

1
2 % Compare the following virtual user interface layouts and select the one
  that best meets the users' requirements. Consider the following factors in
  your evaluation:
3
4 % Ease of Interaction: How easily can users interact with the virtual
  widgets in each layout? Take into account widget positioning, reachability,
  and ergonomics.
5
6 % Visibility and Observation: How effectively are the virtual widgets
  positioned within the user's field of view? Consider both primary and
  peripheral visibility.
7
8 % Spatial Arrangement and Alignment: Evaluate the overall structure and
  alignment of the widgets. Are they organized in a way that is logical and
  efficient for the user?
9
10 % Suitability for Physical Anchoring: Assess how well each layout allows
  virtual widgets to be anchored to specific physical areas or objects as
  required by the user.
11
12 % Overlay and Occlusion Considerations: Determine if any physical areas are
  inappropriately occluded or overlaid by virtual widgets, based on the
  user's instructions.
13
14 % Customization and Adaptability: How well does each layout allow for
  customization or scaling based on user preferences?
15
16 % After evaluating these aspects, select the layout that aligns most
  closely with the users' stated requirements and explain your reasoning.
17 The user's instructions are: [user's instructions]
```



(a) The user's instruction: "I would like the map to be displayed in front of me, with all other applications positioned around it. They should not obstruct the screen."



(d) The user's instruction: "Display the map on the TV, and arrange the flight booking and calculator on the desk so that I can easily interact with them."



(b) The user's instruction: "Could you please show everything to me? I want to reply an email first so place it in the center of the monitor."



(e) The user's instruction: "I need to reply to an email. Please put the email on the monitor and position my calendar and messenger in mid-air next to it. Also, be careful about my coffee mug."



(c) The user's instruction: "I want to buy a Formula 1 car model on Amazon. Can you also display the sky news? Please do not obstruct my monitor."



(f) The user's instruction: "I'd like to watch videos on the TV while scrolling the sports news. I might also play music at the same time, so please position Spotify next to YouTube for easy access."

Figure 12: Instructions and generated layouts in the evaluation study.

D Generated UI layouts in the evaluation

E Module-Wise Time Cost of *AutoOptimization* in the Evaluation

Table 1: Time cost (in seconds) of each module in *AutoOptimization* during the evaluation study. User input time for instructions is excluded. When the ambiguity detection module asks users to clarify, multiple round ambiguity detection times are averaged. Results are presented as mean (standard deviation).

Scenario	Ambiguity Detection	Configuration	Optimization	Validation
Living room	5.11(2.34)	4.54(1.34)	7.85(1.87)	2.12(0.35)
Office	6.18(2.17)	4.35(1.12)	8.46(2.13)	2.34(0.39)
Airplane	5.68(1.96)	5.03(1.37)	6.38(1.29)	2.21(0.41)